


CONNEXIONS®



The Interoperability Report

June 1993 Volume 7, No. 6

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

A National Information Infrastructure.....	2
Network Management: Status and Challenges.....	11
Making a Network Map.....	18
The IETF and OSI.....	26
Announcements.....	29
Book Review.....	31

From the Editor

Speeches or testimonies are not something we usually carry in this publication, but this month we break that rule and bring you the perspectives of two very prominent "Internet Personalities," Vint Cerf and Marshall Rose.

Vint Cerf testified before the US House of Representatives Committee on Science, Space and Technology Subcommittee on Technology, Environment and Aviation on March 23, 1993. The topic was a *National Information Infrastructure*. His written testimony is our first article this month, and it covers the history, present state, and future of the Internet. With a current estimate of 15 million users, the Internet is well beyond being an academic playground. Vint eloquently outlines what possibilities the technology has brought us, and suggests (with input from the Internet community at large) ways in which the government can help foster further developments to establish an information infrastructure, analogous in many ways to the United States interstate highway system.

Marshall Rose was recently named the *Area Director for Network Management* in the Internet Engineering Task Force (IETF). His experience as an implementor and author gives him a unique perspective from which to analyze the state of network management. Our second article is adapted from a keynote address given at the *Third International Symposium on Integrated Network Management*, held in San Francisco, April 18–23, 1993. Marshall has recently completed the second edition of *The Simple Book*. This book describes version 2 of the Simple Network Management Protocol (SNMPv2), and will be available at INTEROP in a couple of months.

Speaking of INTEROP, you may have noticed that we are now calling it "INTEROP 93 August" (rather than "Fall"), to remind attendees of the dates (and avoid redefining the seasons). Mark your calendar for August 23–27 in San Francisco!

I must apologize for my incorrect statement in the April issue regarding the Gopher to X.500 gateway work. This development was done at the University of *Michigan* (not Minnesota).

Our next article, however, most certainly *is* from Minnesota and describes a method for automatically generating complex network maps. Such maps can be a valuable asset for network managers, designers and planners. The article is by Marshall Midden.

The *Internet Engineering Steering Group* (IESG) of the IETF has recently decided to adjust the area boundaries so as to include most of the current OSI activities with other areas. Erik Huizer outlines these changes in our final article this month.

ConneXions is published monthly by Interop Company, 480 San Antonio Road, Mountain View, CA 94040-1219, USA. +1 415-941-3399. Fax: +1 415-949-1779. Toll-free: 1-800-INTEROP. E-mail: connexions@interop.com. Copyright © 1993 by Interop Company. Quotation with attribution encouraged. *ConneXions—The Interoperability Report* and the *ConneXions* logo are registered trademarks of Interop Company. ISSN 0894-5926

A National Information Infrastructure

by Dr. Vinton G. Cerf,
Vice President, Corporation for National Research Initiatives
and President, Internet Society

Written Testimony for the US House of Representatives Committee on Science, Space and Technology Subcommittee on Technology, Environment and Aviation, March 23, 1993.

Introduction

Mr. Chairman, distinguished members of the subcommittee and guests, my name is Vinton G. Cerf and I am Vice President of the non-profit *Corporation for National Research Initiatives* (CNRI). I also have the honor to serve as President of the *Internet Society* (ISOC), which is a professional society of individuals who are users, developers or operators of the Internet. My remarks today are personal in nature, but they are colored by my past and present professional experiences which form the backdrop against which my opinions and observations have evolved.

I worked on the ARPANET project while a graduate student at UCLA in the early 1970s, helping to develop the protocols used to support communication between the computers (hosts) on the network. The highly successful ARPANET experience with packet switching technology led to additional satellite, mobile radio and local area packet networks, developed under *Advanced Research Projects Agency* (ARPA) sponsorship and, in the case of Ethernet, at the Palo Alto Research Center of the Xerox Corporation. Dr. Robert Kahn, now the president of CNRI, initiated an ARPA internetting research program to explore techniques to connect different packet networks in such a way that the host computers did not have to know anything about the intermediate networks linking them together. Dr. Kahn and I developed the idea of *gateways* and wrote the first specification for the basic TCP/IP protocols now used in the Internet.

The Internet

The idea behind the Internet was the seamless linking of many different kinds of packet switched networks. I came to ARPA in 1976 to manage the Internetting research program and by the time I left ARPA in 1982, the TCP/IP protocols were widely used and the Department of Defense had declared them standards for military use. The Internet has blossomed in the subsequent 10 years, particularly after the *National Science Foundation* (NSF) introduced the *NSFNET* as part of the Internet in the mid-1980s. In 1982, there were about 100 computers on the ARPANET and a few score others were part of the NSF-sponsored *CSNET* which also used the *Telenet* public data network. In 1993 there are over 1.5 million of them. The system links over 10,000 networks in roughly 50 countries. Although it is not known for certain how many users there are, we believe there are well over 5 million. The system is tied into most public and many private electronic messaging services and this expands the population able to exchange e-mail to some 15 million. They include business people, academics, government workers, scientists, engineers, librarians, schoolteachers, astronomers, oceanographers, biologists, historians, reporters, attorneys, homemakers, and secondary school students.

Growth

The system is doubling annually in users, networks, hosts and traffic. In some parts of the Internet, such as the NSFNET backbone, traffic growth rates as high as 15% per month have been measured. Internet is growing faster than any other telecommunications systems ever built, including the telephone network. Today, over half of the networks registered are associated with business users.

Of course, these rates of growth cannot continue indefinitely, but there is reason to expect that the user population will exceed 100 million by 1998.

Perhaps even more important, this federal investment in research has created new industries revolving at first around the hardware and software of Internet technology, and more recently, around network and information services supported by the Internet. The new businesses (such as Sun Microsystems, 3Com and Cisco Systems) have highly positive international trade balances and phenomenal growth, commensurate with the rapid growth of the Internet itself. The growth rate is extremely strong in Europe, South America and the Pacific Rim creating major export markets for the US firms offering Internet products and services.

In 1975, operational management of the ARPANET was transferred to the *Defense Communication Agency* (now the *Defense Information Systems Agency*—DISA). In the mid-80s, the National Science Foundation (NSF), the Department of Energy (DoE), and the National Aeronautics and Space Administration (NASA) joined in supporting the evolution of the Internet and developing and applying its technologies. In addition to developing their own networks (that became integral components of the Internet), these agencies participated in the development and standardization of the Internet protocols (the TCP/IP Protocol Suite) and provided support to the secretariats of the *Internet Architecture Board* (IAB) and *Internet Engineering and Research Task Forces* (IETF and IRTF). This included support for the *Internet Assigned Number Authority* (IANA), document editor (RFC Editor), and Network Information Centers which provide information and assistance to users and deal with Internet network address assignments. ARPA, NSF, DISA, DoE and NASA now make up part of the *Federal Networking Council* which continues to oversee the development of networks used in government-sponsored research and education.

The Internet Society

Formed at the beginning of 1992, the non-profit, professional membership *Internet Society* provides an institutional framework for carrying out a variety of activities intended to foster the continued growth, evolution and application of the Internet. Included in this undertaking is the responsibility for the technical standards used in the Internet. Along with members of the Federal Networking Council, the Internet Society supports the IETF Secretariat. It sponsors conferences and workshops on the Internet and its technology, is establishing liaison relationships with the *International Telecommunication Union* (ITU) and the *International Organization for Standardization* (ISO), works with various United Nations agencies (e.g., UN Development Program) to encourage the acquisition and use of Internet facilities in technologically-emerging countries, and participates in efforts to extend Internet services from university and research library communities to secondary school systems.

The Internet Society does not operate any of the thousands of networks that make up the Internet, but it assists service providers by providing information to prospective users and involves product developers and researchers in the evolution of Internet technical standards. Corporate and individual, professional support for this organization is widespread and international in scope.

**High Performance
Computing and
Communications****A National Information Infrastructure (*continued*)**

The *High Performance Computing Act* (HPC) was signed into law late in 1991. The original impetus for this legislation came from then-Senator and now-Vice President Gore whose vision of information superhighways limned the potential of a computing and communications infrastructure which would permeate and stimulate the government, business and private sectors of the US economy. The promise of a vast new economic engine equal to or larger than the engine sparked by the National Highway Act of 1956 was a powerful incentive for this bill and lies at the heart of the motivation for creating a new National Information Infrastructure.

NREN

One of the key elements of the HPC initiative is its *National Research and Education Network* (NREN) program. Designed to extend the performance envelope of networking into billion bit per second (gigabit) territory and to extend the scope of access to a larger segment of the research and education communities, the effort spawned a major research program on gigabit networking. ARPA and NSF jointly funded an effort, organized by the Corporation for National Research Initiatives, to establish multiple gigabit testbeds across the United States. The program is highly leveraged, involving major contributions from the computing and communications industries as well as several of the national laboratories and major research universities.

An important focus of the gigabit testbed program is to discover by experimentation which technologies and applications are likely to form the core of the high performance communication systems of the future. The deep involvement of industry is intended, in part, to assure that the results take into account the plans and capabilities of the private sector. Such partnerships among government, industry and academic institutions form a bedrock upon which new national infrastructure can be founded.

The vision of the NREN component of the HPC effort begins with the existing US component of the global Internet. Under the NREN program, key parts of the US Internet have been extended to operate at 45 million bits per second (in particular the NSFNET) and procurement of higher speed services by DoE and NASA is in progress. The gigabit testbed program is enabling the early availability of very high speed network technology and the results of the program will help to determine the architecture and technology of even higher capacity services. The NSFNET initiative, which began in 1986, has also led to the creation of dozens of new Internet service providers, including a number of for-profit networks offering unrestricted Internet service to all who desire it.

Another fundamental motivation for the high performance networking component of HPC is the intense investment by the principal interexchange and local exchange telecommunications carriers in the US in the use of optical fiber in their networks. Capable of supporting operation in the billions of bits per second, the optical networks form the strands from which a national gigabit fabric can be woven. Investments by local exchange carriers and cable companies to increase the capacity of the lines reaching business and residential customers make it possible to envision a time when very high capacity services can be supported on an end-to-end basis.

The far-sighted vision of the HPC effort, together with the explosive growth of the Internet and basic communications facilities resulting from private sector initiatives, have set the stage for a dramatic new step in the evolution and convergence of computing and communication: the creation of a *National Information Infrastructure*.

Infrastructure

Information Infrastructure is the "common ground" on which computer-based products and services depend to achieve commonality and interoperability. Included in infrastructure are technical standards and the organizations and procedures through which they are developed; communication services and the physical, human and organizational resources needed to deploy, maintain and operate them; legal and regulatory frameworks which encourage cooperative development of precompetitive technology, foster the protection of computer-accessible intellectual property, the protection of privacy, and support the conduct of electronic commerce; widely available computer software for many hardware and operating system platforms establishing ubiquitous and interoperable computing environments in which applications can be embedded. Infrastructure supplies the raw material out of which limitless applications may be constructed.

Some of the characteristics which mark elements of infrastructure include: ubiquity, expandable capacity, simplicity of use, applicability to many uses and broad affordability. A functioning information infrastructure will lower technical and economic barriers to the introduction of computer-based products and services. It will simplify the discovery and ordering of products and services as well as billing for their use or acquisition. It will also facilitate the day-to-day operation of businesses, government, education, health care and all the myriad activities that rely increasingly on the use of computer and communication technology to accomplish their objectives.

Enabling technology

Infrastructure has an *enabling* character. The highway system enabled the suburban housing boom and convenient, door to door delivery of goods. Of course, it also stimulated the automobile industry and travel. The power generation and distribution system enabled the facile application of fractional horsepower motors and a vast array of other electrical appliances wherever they were needed.

Infrastructure development is almost always preceded by critical inventions which motivate the need for the infrastructure. The light bulb preceded and motivated the need for power generation and distribution. The invention of the internal combustion engine and its application in automobiles motivated the need for better roads, service stations, gasoline refining and distribution. Once the roads were in place, their ubiquity and easy accessibility stimulated the production of a vast array of different vehicles, all designed to conform to certain common constraints (size, height, weight) so as to be usable on most of the roads in the system.

The computer is the automobile of the information infrastructure. Laptops are the sports cars; desktops are the sedans; supercomputers are the formula 1 racing engines; and gigantic mainframe data storage systems are the 18 wheelers. The local access networks form the neighborhood streets; high capacity computer networks are the superhighways; and circuit, cell and packet switching systems form the complex interchanges.

A National Information Infrastructure (*continued*)

Just as vehicles on the road can be filled with an endless variety of people and products performing a multitude of services, software applications fill the empty computing vessels to create the new products and services of the information infrastructure. Communication protocols and standards form the rules of the road. When traffic jams and accidents occur, we call on emergency services to assist. The same may prove true for the information infrastructure when viruses infect the system or other software and/or hardware failures occur; we will need comparable emergency assistance to restore critical services and functions.

Cyberspace

The *Electronic Frontier Foundation* (EFF) speaks of computers and computer networking as a frontier in *Cyberspace*. This is an interesting and apt analogy, given the relative immaturity of both technologies. Despite the apparent sophistication of today's computers, networks and software, their application has barely scratched the surface of the latent possibilities. The notion of frontier raises images of boundaries and limits. But Cyberspace is a virtual place. It is created out of software, making Cyberspace an endlessly expandable environment.

Information is, itself, an infinitely renewable resource to be harvested, shaped, applied and recycled. The products and services which can be built atop the computer and communication infrastructure simply have no logical limits. It is this ceaselessly changing, growing, transmuting information resource which will fuel the economic engine of the information infrastructure.

Information Infrastructure formation

The technical challenges to be overcome in creating a national information infrastructure may only be overshadowed by some of the legal and policy problems. Taking the easier ones, first, it should be apparent that standards for the exchange of a variety of types of information (data) are essential. The value of infrastructure is that providers of two services which must interwork do not have to make bilateral agreements with every partner if appropriate technical standards are developed which enable such interworking. In the case of program (software) interworking, common representations of shared information must be agreed upon so that software developers can be reasonably assured that, if they follow the protocols, their application programs will interwork with each other.

Standards

A variety of high and low-level standards are needed for representation of digital documents; information retrieval queries and responses; remote program interactions; financial or other commercial transactions; privacy, integrity and authenticity preservation; and a plethora of application-specific standards for information interchange. These representations need to include the capability for a wide range of media, including sound and pictures. There are a number of representations available for encoding these various media, but there is not yet widespread agreement on a common set. Consequently, we are still some distance away from a workable information infrastructure.

Applications

The applications that can be supported on a suitable information infrastructure are limited only by imagination and creativity. Examples include health care support (e.g., patient information, prescription databases, digitized X-Rays and MRI scans), remote consultation; education (classrooms without walls, using the information infrastructure to receive instruction, explore digital libraries and work with distant partners), manufacturing, provision of government information, and support for electronic commerce (e.g., order entry, electronic or physical delivery of products, electronic payments, product specifications).

Anecdotes from the 21st Century

An important element of Internet growth is the typical pricing strategy of service providers: flat rates based on the bandwidth of the lines used to access the Internet. Unlike some commercial email and other public data network service providers, Internet service providers have not charged by the "packet." Many believe that this policy has had a major, positive effect on the growth of the network because users had little uncertainty with respect to annual costs for use of the system.

Those of us who have lived with the Internet since its inception have been living in what will be common in the next century. In preparation for this testimony, I sent a brief message out on the Internet to hundreds of thousands of people who make daily use of the network. I asked them to offer their thoughts on points they considered important to make. Within hours, I had thousands of responses, not just from domestic sources but from all over the world. Without the infrastructure of the Internet, such a question would not have been worth asking since the answers would have taken far too long to receive, and I could not have applied available computer cycles to sort and sift the resulting responses. My correspondents were almost uniformly enthusiastic about the prospects for national and global information infrastructure. The following were some of the points they made:

- The Internet Society newsletter is created by correspondents all over the globe who e-mail their stories to the editors in Los Angeles, California and Reston, Virginia. The whole process takes places over a few days, with all the editing taking place on-line. Each issue is available on-line within minutes of completion through a variety of information services on the Internet.
- A professor at the University of Southern Louisiana offered to teach a class on Internet use through e-mail on the Internet. 15,000 people applied to take the class! This is distance-learning with clout!!
- A blind student of Shakespeare asked on the net, "where can I get on-line copies of the plays, it's the only convenient way for me to read them." He uses a text-to-speech and text-to-Braille device. He got back many pointers to on-line archives around the world.
- When President Clinton and Vice President Gore were visiting Silicon Graphics in California's Silicon Valley, the audio and video of the speeches were packetized and multicast on the Internet to hundreds of participating sites. This is an example of the nascent potential in combining all forms of communication in computer-mediated form.
- *Internet Talk Radio* recently made the front page of the *New York Times*—it is another example of the convergence of digital computer communications and mass media.
- When I needed information about the Spratley Islands, I just turned to the *CIA World Fact Book* made available on the Internet by the University of Minnesota.
- A technical problem arose with an application running on an Apple Macintosh. The user sent an e-mail message to several distribution lists and news groups and got back helpful responses, some in minutes, from France, Germany, Italy, Australia, India, Singapore, Canada, England, Norway, United States, Finland,... well, you get the idea. Cyberspace has common interest groups that transcend national boundaries.

A National Information Infrastructure *(continued)*

- The city of Wellington, New Zealand, has a computer on the Internet. It has placed there a wide range of information of interest to potential visitors and tourists, local residents, and Internet explorers. There is strong historical evidence that the rich personal interactions that take place on the Internet contribute to a marked increase in face-to-face meetings requiring travel, so the local government is to be commended for its foresight.

Important things the US Government can do

Offered below is a representative set of comments and suggestions received over the course of a few days from the Internet community. Because of its source, it has an obvious Internet bias to it, but despite that, I think these ideas are worthy of serious consideration.

- Invest in the development of pre-competitive software and technology which is made available to industry for competitive productizing. Historically, universities have developed sample implementations of new Internet software which is then used as the basis for product and service development in industry. Occasionally, industry will sponsor development of freely available software which can be readily distributed throughout the network, creating a kind of mini-infrastructure on which more elaborate, for-profit products and services may be based. In both cases, new businesses are often created to service the market created.
- Foster and facilitate the development of technical information standards through cooperative efforts among industry, academia and government. The procedures of the Internet Engineering Task Force are a model for expeditious and effective development because the standards must be implemented by multiple parties and shown to interoperate before they are eligible for standardization.
- Revisit COCOM and US-specific policy on the application, use, and export of the RSA and DES cryptographic technology. Present policies inhibit the creation of particular aspects of global information infrastructure and, in some cases, US companies are placed at a severe disadvantage relative to competitors. These technologies are key elements [no pun intended] in solving problems of intellectual property protection and management and electronic commerce in an on-line environment.
- Adopt the TCP/IP protocols as coequal with the OSI protocols in the US GOSIP specifications (which describe the profile of protocols that are recommended for use in Government procurements). The TCP/IP protocols are already in wide-spread use within the government, so this change would merely acknowledge reality.
- Move aggressively to support library access to Internet services, with particular attention to rural community access.
- Institute training programs to educate the nation's secondary school teachers and support staff on the use of computer and communication technology in the classroom. Subsidize access where this is necessary. Involve state educational infrastructure in this effort. Review highly successful state-level programs as input to national policy development.

- Stimulate the development of quality software for use in curricula at all levels. Consider programs to develop pre-production software and make it available at no charge, leveraging the creativity of national laboratories, universities and individuals.
- Mandate public, on-line availability of government-produced or sponsored information and allow the private sector to add value and resell it. For example, the White House is providing on-line access to unclassified executive orders and text of speeches by senior administration officials within hours (and sometimes minutes) of their release.
- Foster programs to explore and experiment with the use of information infrastructure to support telecommuting. Not only as an energy-saving, pollution-reducing step, but a major tool for implementing the Americans with Disabilities Act provisions. It was noted that home-employment and suburban satellite offices illustrate that electronic communication infrastructure is approaching the importance of the more concrete [pun intended] traffic highways.
- Make use of the Internet to harvest information from its tens of thousands of public databases as an adjunct to intelligence gathering and analysis by various agencies of the federal government. Make available government unclassified information and analysis via the Internet as a contribution to the community (e.g., *CIA World Fact Book*).
- Get all branches of the government on electronic mail and support the ability to exchange email with the public.
- Encourage the deployment of ISDN services.
- Foster the development of shared scientific databases and collaboration tools which can be used to enhance the utility of research results and provide access to raw as well as analyzed data to support corroborating research.
- Make use of the Internet to build bridges among the scientific, research, academic and educational communities.
- Link the museums of the world on the Internet.
- Avoid the unintentional creation of a gap between information rich and poor. The concern here is that private sector entrepreneurship may conflict with freedom of access to public information. Note that the potential gap problem applies equally as well to individuals and to large and small corporations!
- Position national policy so that the government need not subsidize network service providers. Rather, subsidize users, where this is appropriate. By this means, remove most of the *Appropriate Use Policy* dilemmas from consideration at the network level. It is not technically possible today, using existing capabilities, to distinguish different classes of traffic at the network level. [There were a few people who thought the government should build the National Information Infrastructure but the vast majority who commented on this preferred private sector service provision, albeit under government policies which assure ubiquity of service, full interconnection of all service providers and reasonable costs].
- Find a way to make advertising permissible and useful in the National Information Infrastructure.

For further reading

A National Information Infrastructure (*continued*)

[Ed.: The citations below were not part of the written testimony, but are provided here for readers who wish to pursue some of the topics further].

- [1] Lottor, M., "Internet Growth (1981-1991)," RFC 1296, January, 1992.
- [2] Solensky, F., "The Growing Internet," *ConneXions*, Volume 6, No. 5, May 1992.
- [3] Marine, A., "How Did We Get 727,000 hosts?" *ConneXions*, Volume 6, No. 5, May 1992.
- [4] Crocker, D. "The ROAD to a New IP," *ConneXions*, Volume 6, No. 11, November 1992.
- [5] Goldstein, S. & Michau, C., "Convergence of European and North American Research and Academic Networking," *ConneXions*, Volume 5, No. 4, April 1991.
- [6] Roberts, M., "The NREN Takes Shape," *ConneXions*, Volume 5, No. 6, June 1991.
- [7] Roberts, M., "NREN Bill signed into Law," *ConneXions*, Volume 6, No. 2, February 1992.
- [8] Roberts, M., "NREN Politics Thicken," *ConneXions*, Volume 7, No. 1, January 1993.
- [9] Gross, P., "The Internet Engineering Task Force (IETF)," *ConneXions*, Volume 2, No. 10, October 1988.
- [10] Malkin, G., "Inside an IETF Plenary Meeting," *ConneXions*, Volume 7, No. 1, January 1993.
- [11] Postel, J., "An overview of the Internet Activities Board," *ConneXions*, Volume 1, No. 8, December 1987.
- [12] Cerf, V., "Internet Activities Board," RFC 1160, May 1990.
- [13] Chapin, L. (ed)., "The Internet Standards Process," RFC 1310, March 1992.
- [14] Dern, D., "The ARPANET is Twenty," *ConneXions*, Volume 3, No. 10, October 1989.

Author's Address:

Corporation for National
Research Initiatives
1895 Preston White Drive
Suite 100
Reston, VA 22091
Phone: +1 703-620-8990
Fax: +1 703-620-0913

Dr. VINTON G. CERF holds a BS degree in mathematics from Stanford University and a Ph. D. in Computer Science from UCLA. He participated in the development of the ARPANET host protocols while at UCLA from 1967-1972. He served as a member of the computer science and electrical engineering faculty at Stanford from 1972-1976 where he led a research project which developed the TCP/IP protocol suite. He co-authored the original TCP design document with Robert Kahn.

Cerf joined DARPA in 1976 where he managed the Internet, Packet Communications and Network Security programs, departing in 1982, as principal scientist, to join MCI where he served as vice president of MCI Digital Information Services Company. He led the engineering effort to develop MCI Mail. Cerf left MCI in 1986 to become vice president of the Corporation for National Research Initiatives where he is responsible for projects involving the Internet, electronic mail, and Knowledge Robot research.

Cerf is also a fellow of the IEEE and AAAS. A member of ACM, he serves on the ACM Council and the SIGCOMM Executive Committee. He is a member and former chairman of the Internet Activities Board. He serves on several National Academy panels, the Federal Networking Council Advisory Committee, and several corporate technology advisory boards. He was elected to the *Datamation Hall of Fame* and, with Bob Kahn, shared the *ACM Software System* and the *IEEE Kobayashi Awards* in 1992. He was elected President of the Internet Society in 1992 and received the *EFF Pioneer Award* in 1993. He can be reached as: vcercf@cnri.reston.va.us

Network Management: Status and Challenges

by Marshall T. Rose, Dover Beach Consulting, Inc.

[This article is a companion to a keynote speech given at the *Third International Symposium on Integrated Network Management*, held in San Francisco, April 18–23, 1993. In keeping with the spirit of a keynote, the style is conversational. —Ed.]

“In Hell, sinners get exactly what they ask for.”
—*Internet Proverb (undated)*

“The hours are good... though most of the actual minutes are pretty lousy.”
—*Douglas Adams, “The Hitchhikers Guide to the Galaxy” (1979)*

Introduction

The last five years have seen a lot of interest in so-called “standardized network management.” Some of this interest has resulted in products, and some of these products have resulted in solutions. This reflects the natural evolution of the industry: a technology is defined, implemented, marketed, and undergoes constant incremental revision. In this article, I’d like to see if we can understand what is currently going on in the industry, and perhaps even identify the challenges ahead.

Before beginning, I must acknowledge the other members of the SNMPv2 design team—Jeffrey D. Case, Keith McCloghrie, and Steve L. Waldbusser—for their profound impact on the the field of network management. Certainly they have a pervasive influence in the industry, and all the good ideas in this article, I’ve borrowed from them!

“OSI is a beautiful dream, and TCP/IP is living it.”
—*Einar Stefferud, Network Management Associates (1992)*

“Before Political Correctness there was the concept of Open Systems.”
—*Internet Proverb (undated)*

The danger of Dreams

Before looking at the state of technology in the industry, it is important to understand the industry often has difficulty distinguishing between dreams and reality—at least at first. Dreams sell well. They make excellent dogma and even better ad-copy. In times of great stress, they are particularly seductive. Dreams are the stuff of great marketing opportunity. However, dreams, like controlled substances, may be harmful. Although vision is important to progress, “dreaming” often gets in the way of “doing.” In our case, the computer-communications industry spends a lot more time dreaming than coding.

As Stefferud masterfully reminds us, while everyone supports the concept of “open systems,” it takes engineering discipline and numerous production-quality implementations in order to realize a competitive, robust open systems market.

Ultimately however, the industry is able to distinguish between dreams and reality. For example, the “US GOSIP dream” has garnered a lot of hype, many problems, few products, and little credibility. US GOSIP is “checklist procurement” at its finest: a US federal agency specifies GOSIP and TCP/IP, makes sure that each bid has a GOSIP component (possibly to be supplied later), and then proceeds to make technical evaluations based on the TCP/IP products offered. Of course, including GOSIP components (which will probably never be taken out of shrink-wrap) does drive up the price, but that’s what happens when dreams meet reality.

continued on next page

Management: Status and Challenges (*continued*)

As we might expect, the network management area is not immune from dreaming—far from it. Originally, it was the “CMIP/CMOT/CMOL dream.” CMIP (actually the whole OSI experience) has proven quite convincingly that “official” standards bodies are much better about producing international agreements than working technology. Today, it is the “DME dream.” How much longer, I wonder, until people start musing if DME is the OSI of the 90s? Well, when that dream turns into a nightmare, remember you read it here first.

Let’s end our excursion from the Twilight Zone and return to reality.

“The problems of the real world are remarkably resilient to administrative fiat.”
—*Internet Proverb (undated)*

“Young men make wars, and the virtues of war are the virtues of young men: courage and hope for the future. Then old men make the peace, and the vices of peace are the vices of old men: mistrust and caution.”
—*Lawrence of Arabia, Columbia Pictures (1962)*

Management protocol

In the late 80s, we had the network management protocol wars, and when the dust settled, the industry adopted SNMP. A lot of people wonder why the protocol wars were necessary, or regardless, whether the industry can support two standard network management protocols. The answer is that the term “protocol wars” is very misleading. It is not a question of SNMP versus CMIP—it is a question of competing *paradigms* for network management.

At the heart of matter are issues dealing with ease of scaling and behavior during times of network instability. The SNMP camp has always believed that network management must tend toward ubiquity; that is, the addition of management technology must have a minimal impact in order to achieve the widest possible deployment. Further, the SNMP camp has always believed that networks are inherently unreliable, and, as such, the design of the network management framework must be ever mindful of this. So, in a nutshell, the SNMP philosophy is that the network management technology must be widely-implemented; must scale well; must not interfere with normal operations; and, must operate during times of network stress. (It is left as an exercise for the reader as to what principles, if any, are held by the CMIP camp. Certainly, their design choices seem to ignore any of these considerations, e.g., the choice of using a connection-oriented transport is provably fatal for management operations, and yet that’s exactly what CMIP uses!)

When viewed from this perspective, it is clear that the “protocol wars” were inevitable. It should also be clear that any effort to achieve coexistence between the SNMP and CMIP frameworks is doomed to failure. Coexistence—whether through an API or a proxy device—is tantamount to achieving a half a paradigm shift. It simply isn’t possible to find a worthwhile compromise. Indeed, this explains why such efforts have always failed in the past: any API that tries to be “management protocol independent” will capture the worst of both worlds; similarly, proxy mapping schemes between SNMP and CMIP suffer monstrous space-time complexity! These efforts at compromise, whilst politically correct, are fatally flawed in the technical sense.

Periodically, the protocol wars flare up again. Recently, some CMIP proponents have argued that network management requires object-oriented mechanisms which CMIP has and SNMP doesn’t.

Such an assertion lacks coherence. SNMP is based on a remote debugging paradigm, taking the best parts of object-oriented concepts (e.g., defining types or classes and identifying instances of those types, along with support for method routines), but leaves out the questionable aspects (e.g., inheritance and containment hierarchies). The best parts allow for straight-forward description of management information in an implementation-independent fashion, in an efficient and cost-effective manner. In contrast, other parts, included in CMIP, are of questionable value in network management because of the constantly evolving nature of networking technology and network topologies, where hierarchies become out-of-date with each reorganization of the network and each new innovative product on the market.

“No matter where you go, there you are.”

—*The Adventures of Buckaroo Banzai, Twentieth Century Fox (1984)*

“Once again we’ve saved civilization as we know it.
The good news is that they’re not going to press charges.”

—*Star Trek VI, Paramount Pictures (1992)*

SNMP evolution

During a period of over four years of stability and growth, SNMP provided a technology base for a lot of network management products, and a lot of things were learned in the process. Last year, the *Internet Engineering Steering Group* (IESG), the body that oversees standardization of the Internet suite of protocols, issued a request for proposals to evolve the SNMP framework. In its call, the IESG observed that the existing framework provides stable and effective network management for the Internet, which is used pervasively and continuously. As a consequence, any change to the SNMP framework must be evolutionary. Further, the resulting technology must be able to well-coexist with the installed base.

SNMPv2 was developed using a “design team” model, in which four engineers developed a proposal and a working group then took “ownership” to produce the final specification. The entire process from IESG call to initial standardization took a little over a year. Although SNMPv2 is a worthy successor to the original SNMP, the only true test of success is implemented code and wide deployment. Over the next months, we should see whether the SNMPv2’s new features such as security, bulk retrieval, and manager-to-manager interactions, are able to win the hearts and (more importantly) pocketbooks of the industry.

“If the only tool you have is a chainsaw, then everything starts to look like a tree.”

—*Internet Proverb (undated)*

“If network management is viewed as an essential aspect of an internet, then it must be universally deployed on the largest possible collection of devices in the network.”

—*Corollary to “The Fundamental Axiom of Network Management”*

Management Objects, Agents, and Desktops

In the SNMP framework, groups of related management objects are collected into a *Management Information Base* (MIB). Historically, virtually all standardization has dealt with the wire, but now this is moving to systems and applications management.

Management: Status and Challenges (*continued*)

The philosophy for defining management objects has been to provide both standard and vendor-specific definitions. The standard definitions provide for the core functionality associated with a technology, whilst each vendor's definitions are specific to a given product. At present, there are over 2000 standardized objects. In addition, a few hundred vendors have written their own proprietary object definitions. (Rumor has it that one router vendor has defined nearly 2000 objects for just one product.)

Today, SNMP agents run on just about everything, from intermediate-systems, such as routers and gateways; supporting devices, such as bridges, repeaters, hubs, and modems; and end-systems, such as workstations, mainframes, terminal servers, and printers. However, there is still an area where we need more work, network management of the desktop. This is an important area as it forms the basis for harmonizing desktop and internet management in the enterprise.

Desktop computing can be characterized as a hardware/software platform which supports numerous third-party components. The trick to enabling desktop management is to make it easy for the third-party vendors to add network management instrumentation to their components. There is a lot of interest in this area, and several activities. Hopefully, we will soon know how long it will take to enable the desktop for network management.

"When the going gets weird, the weird turn pro."

—Hunter S. Thompson, *"Fear and Loathing in Las Vegas"* (1971)

"In (the Soviet Union), they have it all mapped out so that everyone pulls for everyone else. But what I know about is Texas, and down here you're on your own."

—*Blood Simple*, Circle Films (1983)

Management platforms

In the SNMP approach, the complexity of network management resides with platforms and applications. As might be expected, this is the area where the least amount of progress has been made. One reason for has been that the industry's initial approach was to develop graphical platforms for management applications, rather than developing the actual management applications.

Ideally, a management platform provides "integration" services. That is it deals with information presentation (barometers and gauges); information control (buttons and menus); and, mapping data to information. Unfortunately, today's management platforms have dealt primarily with the first issue, a little with the second issue, and have done virtually nothing in the area of turning data into information.

As such, today's management applications are on their own. The result, hasn't been a pretty picture.

"They're as rare as hen's teeth."

—Dr. SNMP, quoted in *"The Simple Times"* (1992)

"Everyone wants results, but no one is willing to do what it takes to get them."

—*Sudden Impact*, Warner Bros. (1983)

Management applications

Thus far, we've seen three kinds of applications: browsers; mappers; and, (very few) thinkers. Although browsers aren't "real" management applications, it's amazing how quickly a well-trained operator can solve problems by simply perusing remote data.

Mappers (so-called auto-discovery tools) “clean up and show well,” but are probably no substitute for a little bit of facilities discipline. But where are the thinkers?

There are two reasons: first, management objects are chosen from a cross-section of the available data, rather than with an eye towards solving specific network management problems. That is, management objects are defined for generic network management, not as a means for solving a specific problem. Second, it takes a large amount of diverse information to describe an operational network and (network) management practices differ quite a bit between and within organizations. In other words, few organizations are run the same way—each has different policies. So, should we really expect third-party vendors to be able to produce management applications that will run in a variety of actual networks?

This suggests that we need two things: management objects that are problem solvers; and, a technology for relating devices within a network. Before looking at a framework to achieve this, let's look at the last part of the state of the technology.

“This is precisely the sort of thing that no one ever believes.”

—*The Adventures of Baron Munchausen*, Columbia Pictures (1989)

Dual-role entities

Recently, there's been a fair bit of activity in so-called “dual-role entities,” devices that are both agents and managers. These devices collect and process information from agents and then make this information available to managers. Examples include devices that implement SNMPv2's manager-to-manager facilities, RMON (remote network monitoring) probes, and proxy devices.

This technology offers a lot of intriguing possibilities. A dual-role entity could be used support delegated management, in order to allow the network management system to better scale. As a part of this, a dual-role entity could synthesize new data based on the data it has collected. That is, instead of simply filtering, a dual-role entity might process the data in order to derive new information.

“Spread the manure where you are going to grow the vegetables.”

—*Dr. SNMP*, quoted in *The Simple Times* (1992)

“There can be only one.”

—*Highlander*, Twentieth Century Fox (1986)

Manager framework

So, what does the industry need? Perhaps it's time to focus on a framework for management applications, and to realize this by providing a reference management platform.

To begin, it should provide a “standard but minimal” API, which hides the details of protocol versions (SNMPv1 or SNMPv2); transport backing (UDP, IPX, etc.); and, application-level addressing (SNMPv2's parties, contexts, etc.) At present, each brand of management station has its own API, which has an inhibiting effect on applications development. Fortunately, it is possible to build an API which transparently supports both SNMPv1 and SNMPv2. The reason is that SNMPv2 is an evolution of, and is wholly consistent with, the original SNMP framework. (As noted earlier, it isn't possible to build such an API spanning both SNMP and CMIP because their underlying frameworks are too different.)

Management: Status and Challenges (*continued*)

Beneath this API, we need an engine to relate MIB modules (properties); devices (capabilities); and, components (instances). For example, a particular router product can be viewed as having a number of capabilities, defined in terms of the management objects it implements from several MIB modules. Each router in the network is a component with an instance of this information. Of course, a component could be something other than a physical device, e.g., a “network cloud” service, such as ATM, Frame Relay, or SMDS.

Once we have this API in place, we need to develop a framework to define “knowledge templates” to relate object assertions and the actions which should occur when an assertion fails. This would then provide a basis for “problem solving” management objects. Of course, such a framework couldn’t be developed overnight, and a fair bit of experimentation is needed to get the right thrust/payload ratio.

Note that these “object-oriented” concepts are realized solely at the management station, and are unnecessary, either at the agent, or between the management station and the agent. This is an important observation for two reasons: first, it doesn’t unnecessarily burden the agent. The goal of an agent is to efficiently export management instrumentation—it is unreasonable to expect a device to exhibit allomorphic behavior in order to be consistent with an object-oriented paradigm. Second, by realizing these concepts at the management station, we are free to experiment with the framework, to refine our understanding, and perhaps even to learn something.

Of course, in order to ease adoption, such a framework probably requires both a vendor-neutral reference implementation and an applications repository. Both must be openly-available. Of course, with these in place, a “cottage industry” of third-party management application developers would be enabled, and then perhaps we’ll start to see some real management applications.

“A solution encompasses a product. A product embraces a technology.
A technology implements a standard. They’re different.”

—Michael D. Zisman, *Soft•Switch* (1991)

“Why are Communists preferable to Standards Committees?
Because Communists make 5 year plans.”

—Paul V. Mockapetris, quoted in *“The Little Black Book”* (1991)

Conclusions

“Standardized network management” has certainly had a great impact on the industry. But, it is important to distinguish between two kinds of standards, the horizontal and the vertical. Horizontal standards, such as management protocols and management objects, can, at most, provide a technology base for enabling products. However, in order to produce a robust market, vertical standards, such as APIs and a framework for management applications, are needed.

Meanwhile, the industry continues its struggle to distinguish between dreams and reality—“architects” and marketers are very good at seducing the industry’s attention for long periods at a time. Fortunately, it is the engineers who ultimately must deliver products which provide cost-effective solutions. As such, the outcome is certain. The only thing uncertain is how long it will take until the inevitable is acknowledged.

References

- [1] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Introduction to version 2 of the Internet-standard Network Management Framework," RFC 1441, April 1993.
- [2] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1442, April 1993.
- [3] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Textual Conventions for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1443, April 1993.
- [4] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Conformance Statements for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1444, April 1993.
- [5] Galvin, J. M., & McCloghrie, K., "Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1445, April 1993.
- [6] Galvin, J. M., & McCloghrie, K., "Security Protocols for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1446, April 1993.
- [7] McCloghrie, K. & Galvin, J. M., "Party MIB for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1447, April 1993.
- [8] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1448, April 1993.
- [9] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Transport Mappings for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1449, April 1993.
- [10] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1450, April 1993.
- [11] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Manager-to-Manager Management Information Base," RFC 1451, April 1993.
- [12] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Coexistence between version 1 and version 2 of the Internet-standard Network Management Framework," RFC 1452, April 1993.
- [13] *ConneXions*, Two Special Issues on Network Management and Network Security, Vol. 3, No. 3, March 1989 and Vol. 4, No. 8, August 1990.
- [14] Case, J. D., Davin, J. R., Fedor, M. S., & Schoffstall, M. L., "Network Management and the Design of SNMP," *ConneXions*, Volume 3, No. 3, March, 1989.
- [15] Case, J., McCloghrie, K., Rose, M. T., Waldbusser, S., "The Simple Management Protocol and Framework: Managing the Evolution of SNMP," *ConneXions*, Volume 6, No. 10, October 1992.
- [16] Rose, M. T., "Network Management is Simple: You just need the 'Right' Framework." In *Integrated Network Management, II*, Iyengar Krishan and Wolfgang Zimmer, editors, pages 9-25, North Holland, April 1991.

MARSHALL T. ROSE is Principal at Dover Beach Consulting, Inc., a California-based computer-communications consultancy. He spends half of his time working with clients, and the other half involved in self-supported, openly-available projects. Rose lives with internetworking technologies, such as TCP/IP, OSI, network management, and directory services, as a theorist, implementor, and agent provocateur. He is the author of four professional texts—on Open Systems Interconnection (*The Open Book*), internet Management (*The Simple Book*), OSI Directory Services (*The Little Black Book*), and Electronic Mail (*The Internet Message*)—all published by Prentice-Hall. Rose received the Ph.D. degree in Information and Computer Science from the University of California, Irvine, in 1984. His subscriptions to *The Atlantic* and *Rolling Stone Magazine* are in good standing. E-mail: mrose@dbc.mtview.ca.us

How to create a Network Map

by Marshall Midden, University of Minnesota

Introduction

There are many reasons for desiring a current network map. There is an inherent complexity resulting with the volume of information that is necessary to make a map of a large network. This article attempts to answer the familiar questions: How, Who, What, Where, When, and Why in reverse order. The questions will be answered as they pertain to the University of Minnesota, as different criteria will lead to different results.

Why create a network map?

There are many reasons for creating a network map. It's "neat and cool." You can show your boss what you do. It is documentation of the present state of the network. It is useful for recording history. Problem diagnosing can be made easier. Statistics and usage data may be interpreted. Growth and planning of a network requires an understanding of what equipment existed in the past as well as what exists in the present.

Where and when?

A person should be able to print it locally, and it should show all five campuses of the University of Minnesota, and various geographically isolated sites. It should be able to be posted on a wall of the office.

The map should be somewhat easy to update, and be done about once per month. A map should be able to be put together by someone else in the University, if they should happen to want one.

What should be on the map?

All active networking equipment should be on the map. Information should include all routing information and all networks. The University of Minnesota routes TCP/IP, AppleTalk Phase II, DECnet, and IPX. There are a couple of OSI routes, but no one really uses OSI and it is ignored for now, but the ability to add additional network information must be provided.

Who are the maps for?

Who is going to use it? Network staff (co-workers) will use it for problem diagnostics. The department director will use it for promoting the department. Other people around the campus will use it for reference. It will be used in combination with statistics for growth and planning.

Who is going to create it? Me. I did an earlier map several years ago. I pawned it off on someone else the next time, and now everyone knows how difficult it is. (More later.)

Who is going to put it together? You might be wondering, "how big is this map going to be?" The last hand created map was one page—granted a complicated page with extremely tiny writing, but it was still one page.

Calculations

Let's perform some calculations using the fixed width Courier font at a point size of 11. Lets assume that a device is symmetric in height and width when printed. Routers are printed as circles, for example. Names of devices can be up to about 30 characters long (telecomm-fastpath-1.gw.umn.edu). One point is 1/72 inches. The width of a fixed point Courier font is 0.6 times the height.

$$\text{Area is } \left\{ \left(\frac{30 \text{ characters}}{\text{name}} \right) \times \left(\frac{11 \text{ points height}}{\text{character}} \right) \times \left(\frac{0.6 \text{ points width}}{\text{point height}} \right) \times \left(\frac{1 \text{ inch}}{72 \text{ points width}} \right) \right\}^2 = \left\{ \frac{2.75 \text{ inches}}{\text{name}} \right\}^2$$

The area needed for a device then is 2.75 inches by 2.75 inches. A matrix of 3 devices by 4 devices can be placed on an 8.5 by 11 inch sheet of paper. Using a tree structure uses more space, as does putting in network lines. If we ignore the Shiva Fastpaths, 10Base-T hubs, bridges, etc., and only look at the 78 routers, the minimum map size possible, with no network lines, would be 3 pages by 3 pages.

The current University of Minnesota map showing the routers is 7 pages across by 4 pages high. This shows the logical layout of routers, AppleTalk Numbers and Zones, Novell network numbers, primary and secondary IP addresses of the interfaces, type of device, interface names, DECnet numbers, etc. With the font reduced to 3 points (the smallest almost readable font at 300 dots per inch), we can show the 260 Fastpaths and 280 10Base-T hubs, and keep the size at 7 pages by 5 pages. If the 14,000 other IP devices were added—do you have a mall and stock in 3M?

Gathering the information

About three years ago, a map was created on a Macintosh. There were ten routers, and it was easy. The location of each Fastpath was pretty much unknown, no 10Base-T hubs existed, and life appears to have been simple. The second version had all known internet subnets displayed, and the text got tiny. A project was started to stabilize the AppleTalk networks. One outcome was the creation of a file containing all the AppleTalk gateways.

A backbone redesign effort saw the purchase and installation of many routers. There was no way to fit a map on one page. Things were changing so fast that it was difficult to keep up with what was where. Our department mostly controls the DNS (*Domain Name System*) for `umn.edu`. We insist that all devices have the IP addresses of all interfaces listed in the DNS. How do you enforce this? Anyone can grab an IP address and start using it. A program was written that grabs router information nightly, including configurations and ARP tables. In addition, all active subnets have all IP addresses *pinged* once a month. Scripts are then run to determine differences and potential problems.

Any active device that has other devices behind it, and is put on the University's network is recommended to have MIB 1 SNMP (*Simple Network Management Protocol*) completely installed, with all write options disabled for security. There is a small list of devices that we do not control, and when it is map generation time, the interfaces and speeds are collected using SNMP tools.

There are devices we know about but are not able to do SNMP, but have DNS entries that we control (and thus we can put comments with the DNS entry). This list is about five times the size of the SNMP list, mainly due to a remote campus controlling their DNS entries.

Verifying data

Managing and verifying that data is correct sounds much easier than it is. As the data is continually changing, and more than one person is responsible for it, the possibility of bad data quickly increases. Input for the map generation program comes from at least five different sources:

- *AppleTalk gateway file* (288 entries): This information is checked nightly with SNMP. This file is corrected and updated daily. The chance of this data being correct is good.
- *Domain Name System* (DNS): comments with important devices. The DNS is changed frequently—many times per day. We are using comments to highlight devices/addresses/hosts that are worth monitoring. This data is changed by several people. The expectation of bad data is high due to syntax problems, unreasonable data, impossible situations, etc. Comparing the AppleTalk information file with the DNS entries finds one or two errors per month in the AppleTalk file, and double that in the DNS (usually a new device that was not entered into the DNS correctly).

How to create a Network Map (*continued*)

- *Router configurations:* 55 Cisco routers. The router configurations are gathered every night from the routers. In addition, this information tells us about static routers—like Suns, Novell servers, etc. The map generating program validates the routes between all the routers. This program usually finds one problem per month, usually extremely minor. The program then validates with the DNS information and the AppleTalk information. Any interfaces or IP addresses missing in the DNS or routers, overlapping of AppleTalk gateways, wrong types of devices, or unreachable devices are reported as problems.
- *Static information:* hand entered. This static information is verified with the DNS, AppleTalk, and router configurations. Any data that is suspect is reported. This information tends to go stale, due to not being used for any other purposes.
- *SNMP information:* seven entries. Many serial lines on routers do not have their line speeds entered into the router configurations.

Many pages of warnings are issued as overlapping and replicated data occurs, line speeds and device types are updated and validated, etc. With luck, few real errors are found.

The UM monitoring and mapping hierarchy

Our network is designed as a star network (Ethernet fiber rings have been eliminated). A set of core routers are linked together with an FDDI ring. We think of the backbone network as existing in the backplanes of the routers. These core routers are located in the same rack in our Telecommunications building.

There exists fiber between a total of about 60 buildings located on the: St. Paul campus, East and West Banks of the Minneapolis campus, and a few nearby buildings. This is 50 micron fiber installed in 1986. These buildings get their network connectivity from individual interfaces on a Telecommunications core router. The University's IBX phone switch provides a service known as *LANmark*, which can be thought of as a one megabit Ethernet connection. (A LANmark group is a network. Each LANmark group has its traffic isolated from other LANmark groups.) Separate LANmark groups are used to connect buildings that do not have fiber, to their own ports on the core Telecommunications routers.

An attempt to limit the number of router hops has been imposed on the design. Our remote campuses and large users may have a router that has other routers behind them. There does exist a couple of redundant paths and load sharing links. These are noted, but ignored in map creation.

I must digress at this point into network monitoring. Our custom made network monitoring system is hierarchical in nature. This means that if something in the middle of path to a device that is being monitored goes down, all devices beyond that device are no longer monitored, until the device that has gone down returns. This dramatically reduces network load when a problem occurs. By not pounding on a device that potentially is having problems due to saturation, the network becomes more stable. The network map generation program was originally written to generate the configuration file for our custom network monitoring program. To verify that the 832 devices being monitored (4176 IP addresses for SNMP usage information) were correctly being monitored in the hierarchy, an option to generate a map was added (OK—several dozen options and parameters).

The hierarchy is started from a supplied IP address or network and netmask (assumed to be an eight bit subnetted network). Then using all the supplied data, the program finds everything on that network and marks them as on that network. Recursively going through all interfaces on all marked devices, the second level networks are found, then third, etc. All previously reached devices are noted as either having multiple links or redundant paths. At the end, all devices left over are printed out as unreachable and in error. This happens when a department moves but the DNS entries are not deleted or updated, or a bad IP address was encountered someplace.

With the hierarchy determined, there remains two parts to printing: 1) the sizing and placement of objects, and 2) the creation of the printing output. The sizing is separated because it is not readily determined the best way to display the network diagram. Where to display the information for each device and network interface was determined by the margins of life's experiences. Devices with multiple interfaces have the interfaces positioned equidistant around a circle. See below.

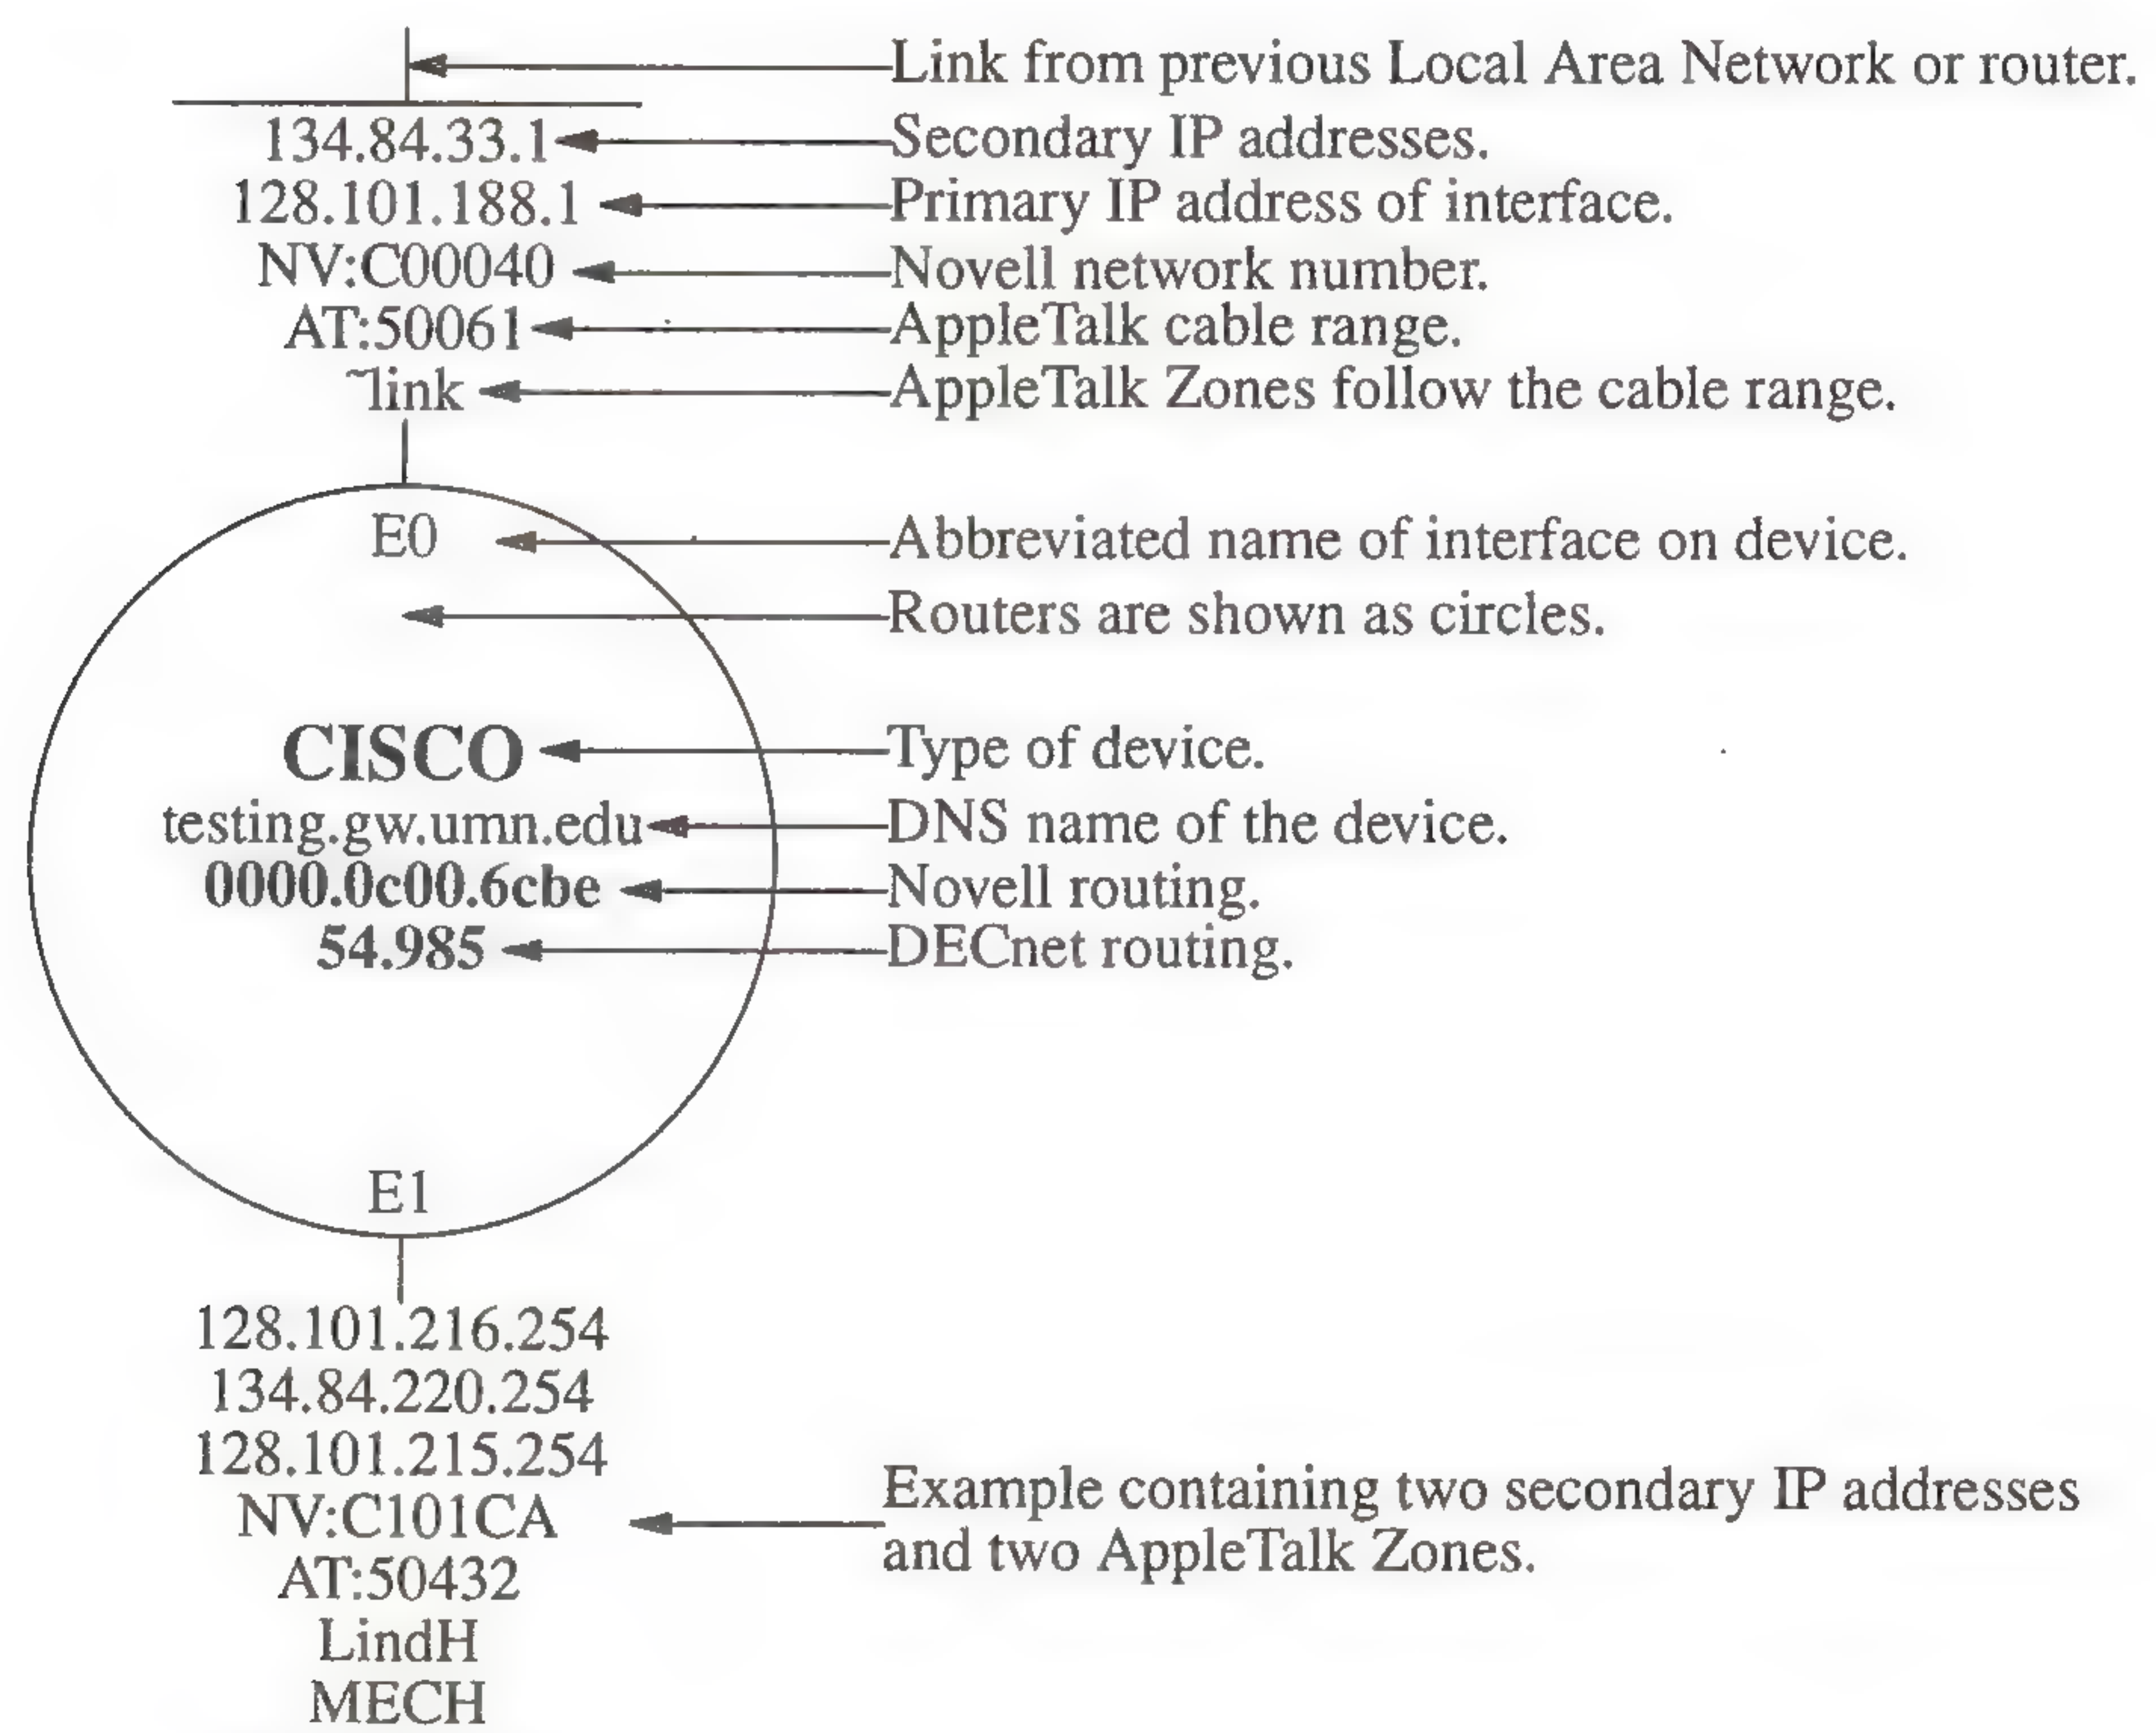


Figure 1: Key to map entry

Position placement of map objects

The algorithm used for placement of devices is the least area of non-overlapping rectangles. In other words:

- If a leg of a router overlaps any other leg, extend that leg outward a tad;
- Go to the next leg and possibly extend it.
- Continue till no leg overlaps any others.
- Next shorten the first leg as much as possible, starting from the shortest possible length, and by small increments extend it outward until it no longer overlaps.
- Continue shortening till all legs have been shortened.

How to create a Network Map (*continued*)

This process creates a reasonable looking map, without the necessity of attempting to calculate the smallest non-overlapping area with geometry and trigonometry. In other words my math was not up to it, and I was under a self-imposed deadline.

Doing this immediately showed that one needed to change the position of router interfaces so that a smaller area would be occupied by the router and its sub-networks. When every possible interface position was tried, and the smallest area taken, computational time limitations were reached (i.e., it took longer than a weekend to finish). When the output of smaller networks was examined, I noticed that typically the smallest total area occurred when the larger interfaces of a router were positioned opposite the interface that led to this router. This lead to a simple sort by area of all router interfaces, and a bell curve approach used. After this sort, then the lengthening and shortening of the interfaces is done, as explained above.

After a router has had all interfaces positioned, it is placed onto a network, just to the right of the last router, if any. No optimization of placement is done here, although it is on the desired features list.

Optimization of the area overlapping algorithm has been done over time. There is now an outer rectangular area that contains all items on a router leg. If this doesn't overlap, none of the individual areas (router, each text string, LAN lines, etc.) are checked. If it does overlap, then both areas must have all areas within checked for overlap. This is a computational intensive effort, and some optimization of the routines involved has been done leading to obscure code.

Printing the map

Printing presented its own challenges. A problem that appeared was due to sorting after an area was laid out. You can not lay something out, then move its orientation and expect text to remain symmetric, especially if you are not rotating the text. All movements due to sorting are noted, and the orientation layout is repeated until no sort changes occur, or a defined limit for those cases that toggle. Just before printing, sorting is turned off and a final positioning is done.

Finally the layout can be printed. I started by generating *pic* (*[gnt]roff*), because *pic* automatically sizes to keep things on one page. Then I switched to Calcomp and used a Calcomp to *PostScript* converter to print it. Quickly it became obvious that if I was printing *PostScript*, then *PostScript* should be generated. *PostScript* need not be complicated. If one ignores definitions and does no computation, conditionals, or looping it is much simpler. Circles, lines, text placement, and the like are easy. Due to limitations on program size for looping, multiple pages need to be clipped before printing. At the time of clipping, it was easy to put in cutting lines, page numbers, etc.

Samples

Figure 2 shows the University of Minnesota routers on one page. Notice that the map, at this size, is unreadable. Expanding it to 7 pages by 5 pages allows for clarity. Figure 3 shows the 10Base-T hubs, Fastpaths and routers, without any text. Putting everything considered important enough to be monitored makes things so small as to be worthless, or so large as to require a mall, depending on the aspect taken.

Map on March 8, 1993 @ 14:23:20 CST




Figure 2: University of Minnesota Routers

How to create a Network Map (continued)

Map on March 8, 1993 @ 14:23:20 CST




Figure 3: 10Base-T hubs, Fastpaths and Routers

References

- [1] Marshall M. Midden, "Experiences in Networking at the University of Minnesota," pp. 36–45 in Proceedings of SUUG November 1991, SUUG, Moscow, USSR (November 1991).
- [2] Craig A. Finseth, "Thoughts on Network Management at the University of Minnesota," *ConneXions*, Volume 6, No. 12, December 1992.
- [3] Craig A. Finseth, "SLIP at the University of Minnesota," *ConneXions*, Volume 7, No. 1, January 1993.
- [4] Case, J. D., Fedor, M., Schoffstall, M. L., Davin, C., "Simple Network Management Protocol (SNMP)," RFC 1157, May 1990.
- [5] McCloghrie, K., Rose, M. T., "Management Information Base for network management of TCP/IP-based internets," RFC 1156, May 1990.
- [6] Rose, M. T., McCloghrie, K., "Structure and identification of management information for TCP/IP-based internets," RFC 1155, May 1990.
- [7] McCloghrie, K., Davin, J., Galvin, J., "Definitions of Managed Objects for Administration of SNMP Parties," RFC 1353, July 1992.
- [8] Galvin, J., McCloghrie, K., Davin, J., "SNMP Security Protocols," RFC 1352, July 1992.
- [9] Davin, J., Galvin, J., McCloghrie, K., "SNMP Administrative Model," RFC 1351, July 1992.
- [10] *ConneXions*, Two *Special Issues* on Network Management and Network Security, Vol. 3, No. 3, March 1989 and Vol. 4, No. 8, August 1990.
- [11] *ConneXions*, Two *Special Issues* on Internet Routing, Vol. 3, No. 8, August 1989 and Vol. 5, No. 1, January 1991.
- [12] Marshall T. Rose, *The Simple Book—An Introduction to Management of TCP/IP-based internets*, Prentice-Hall, 1990, ISBN 0-13-812611-9.
- [13] Jeffrey D. Case, James R. Davin, Mark S. Fedor, & Martin L. Schoffstall, "Network Management and the Design of SNMP," *ConneXions*, Volume 3, No. 3, March, 1989.

(Various maps and the source for this program, "p," are available via anonymous FTP on mail.unet.umn.edu, under the directory unet/maps.)

MARSHALL M. MIDDEN has been programming since 1969. Since 1981 he has worked for the University of Minnesota, on mainframes, minis, micros, supercomputers, and in 1989 was the first employee of a new department called University Networking Services. He has a tendency to get involved with a variety of tasks, from network management, compiler optimization, application development, real-time graphics, military subcontracts, to evaluating Uninterruptible Power Supplies for installation on the University Network in asbestos filled basements. Marshall's e-mail address is: m4@umn.edu

The IETF integrates OSI related work

by Erik Huizer, SURFnet

Introduction

Global internetworking is in a turmoil. What is normal today seems old and bygone tomorrow. Lively OSI debates continue outside as well as inside the Internet community [1–8]. The *Internet Engineering Task Force* (IETF) works hard on defining a new version of IP. The organizational structure of the IETF gets formed into one that is on one hand in tune with the belief in “rough consensus and running code” [9] and on the other hand copes with the explosive growth of this international open standards body. ISO makes liaison statements towards the IETF. Novell wants to publish IPX as an informational RFC. There is a steady increase of participation from commercial companies in the IETF [10]. And last but not least the Internet has become *multiprotocol*.

Amidst all this, the “normal” IETF standardization work is still going on [11]. Mostly unperturbed, but of course not unaffected, by these events. Most of the standards work that the IETF does is still related to the TCP/IP protocol suite, however an increasing amount of work is done on multiprotocol. AppleTalk, IPX and SNA are increasingly used on the Internet, forcing the IETF to deal with the multiprotocol issues. Fortunately at the same time the IETF is getting staffed with able people from various organizations to work on these issues.

The OSI suite of protocols as defined by the CCITT and ISO, and mandated by government administrations, is seen by the IETF as an important set of open protocols for integration in a multiprotocol Internet. As a result of this the last four years a lot of work has been ongoing in the IETF on issues relating to OSI [12].

OSI integration

This work is done in IETF *Working Groups* that resided in an area called the “OSI Integration Area.” This area, and the working groups within it, were formed to work on issues related to OSI based protocols, that were specific to the Internet.

The working groups in this area mostly work on:

- Gatewaying (applications) & interworking/bridging (lower layers) issues for related OSI and TCP/IP based protocols; (e.g., MIME to X.400(88) bodypart mapping definition [13])
- Defining a functional standards for deployment of an OSI based protocol on the Internet (e.g., defining how to store a presentation address in the X.500 Directory Service [14], defining *ping* for CLNS [15]).
- Operational aspects of OSI and TCP/IP protocols on the same backbones.

Besides these issues work is being done on “hybrids” (OSI applications running on top of TCP over IP, and TCP/IP applications running on top of TP4 over CLNS [9]) and on “crossovers” (e.g., mapping of SNMP onto CLNS/TP).

Pilots

The work of these working groups has resulted in the deployment of various OSI originated protocols on the Internet, some operational some still only as pilots. These pilots and services provide us with the experience necessary to decide how to integrate (or gateway) and use these services on the Internet. With this work on OSI integration the IETF makes sure that the Internet community can benefit from (parts of) the OSI protocol suite.

Major X.500 pilots are under way on the Internet [16]. Several CLNS pilots are under way, and several networks already deploy CLNS in an operational environment [17]. And though SMTP/RFC 822 mail is still by far the most deployed mail protocol, various parts of the Internet (especially in Europe) are now using X.400 based mail.

All this adds to the Internet being multi-protocol. At the same time this requires a good coordination between IETF working groups from different areas working on the same type of service, but based on different protocols (e.g., SMTP/RFC 822 extensions, MIME, MIME-X.400(88), X.400 operations, remote mail and mail requirements all deal with the E-mail service).

New structure

For most of the time, there have been two co-directors for the OSI integration area, with one of the directors focusing on upper layer issues and the other focusing on lower layer issues. Recently, after reviewing the increased intertwining of OSI activities with the work in the other areas, particularly in the Internet and Applications area, the *Internet Engineering Steering Group* (IESG) decided to adjust the area boundaries to include most of the current OSI activities with these areas. The current OSI directors, Dave Piscitello and Erik Huizer, have joined the Internet and Applications areas respectively. All of the current work in progress is continuing unchanged.

This leaves us with a structure where all working groups involved with the Internet network layer (be it IP, AppleTalk or CLNS) are grouped together in one area, the Internet area. Similarly all working groups dealing with e-mail are now grouped together in the Applications area. The IESG is confident that this reshuffling of the working groups will make the IESG more efficient and more generally able to deal with multiprotocol issues. This will result in a better coordination between related working groups, and lead to a better integration of OSI-based services in the Internet. For a current list of IETF working groups, retrieve the file `lwg-summary.txt` in directory `/ietf` on host `cnri.reston.va.us`

References

- [1] Jacobsen, Ole, "The Trouble with OSI," *ConneXions*, Volume 6, No. 5, May 1992.
- [2] desJardins, Richard, "Opinion: OSI is (Still) a Good Idea," *ConneXions*, Volume 6, No. 6, June 1992.
- [3] Metcalfe, Bob, Smart, Bob, and Blackshaw, Bob, "Letters to the Editor," *ConneXions*, Volume 6, No. 6, June 1992.
- [4] Rose, Marshall, "Comments on: 'Opinion: OSI is (Still) a Good Idea,'" *ConneXions*, Volume 6, No. 8, August 1992.
- [5] desJardins, R., "Comments on: 'Comments on: Opinion: OSI is (Still) a Good Idea,'" *ConneXions*, Vol. 6, No. 10, October 1992.
- [6] desJardins, Richard, "Internet 2000," *ConneXions*, Volume 6, No. 10, October 1992.
- [7] Hoffmann, Harald, "Letter to the Editor," *ConneXions*, Volume 6, No. 11, November 1992, p 31.
- [8] Crocker, D. & Feit, S., "Letters to the Editor," *ConneXions*, Volume 7, No. 2, February 1993, p 23-25.
- [9] Quoting Dave Clark from a talk he gave at the Cambridge IETF meeting in July 1992.
- [10] Readers not aware of these developments should send e-mail to: `ietf-request@cnri.reston.va.us` and ask to be added to the IETF electronic distribution list.

continued on next page

IETF integrates OSI related work (*continued*)

- [11] RFC 1310 documents the IETF standards process. The reader is advised however that due to the developments referred to in the first paragraph of this article, RFC 1310 may soon be replaced.
- [12] Any resemblance between the four years mentioned here and the four year revision cycle that ISO/CCITT maintained until recently is purely coincidental.
- [13] This work is still ongoing, though nearly finished. To tune in to the final discussions send mail to: `mime-mhs-request@surfnet.nl`.
- [14] See RFC 1277.
- [15] See also *ConneXions*, Volume 3, No 10., October 1989. This work is still ongoing, in the Network OSI Operations group. To join send mail to: `noop-request@merit.edu`.
- [16] See for example RFC 1006 and the ISODE (software).
- [17] a.o. Paradise project in Europe, White Pages project in the US, and AARNet Directory Project in Australia (see also *ConneXions*, Volume 3, No. 6, June 1989).
- [18] Examples of networks running CLNS are NEARnet and ESnet. See also: *Computer Networks and ISDN Systems*, 25 (1992) 4–5.

ERIK HUIZER holds a M.Sc. and a Ph. D. from Delft University of Technology. Since 1988 he has worked for SURFnet, the Dutch academic and research network. He is responsible for various projects including "E-mail to the DeskTop," "X.500 Directory Services" and "Multimedia E-mail." He is a member of various RARE and IETF working groups, of the RARE Technical Committee and of the Internet Engineering Steering Group. He can be reached as: `Erik.Huizer@SURFnet.nl`

"Components of OSI" in *ConneXions*

Integrated Services Digital Network (ISDN)	April	1989
X.400 Message Handling System	May	1989
X.500 Directory Services	June	1989
The Transport Layer	July	1989
Routing overview	August	1989
IS-IS Intra-Domain Routing	August	1989
ES-IS Routing	August	1989
The Session Service	September	1989
Connectionless Network Protocol (CLNP)	October	1989
The Presentation Layer	November	1989
A taxonomy of the players	December	1989
The Application Layer Structure	January	1990
File Transfer, Access, and Management (FTAM)	April	1990
The Security Architecture	August	1990
Group Communication	September	1990
X.25—the Network, Data Link, & Physical Layers	December	1990
The Virtual Terminal ASE	January	1991
Systems Management	April	1991
CO/CL Interworking	May	1991
Open/Office Document Architecture (ODA)	August	1991
Abstract Syntax Notation One (ASN.1)	January	1992
Broadband ISDN	April	1992
Synchronous Optical Network (SONET)	April	1992
Asynchronous Transfer Mode (ATM)	April	1992
Inter-Domain Routing Protocol (IDRP)	May	1992
The Remote Procedure Call (RPC) Service	June	1992
OSI Conformance Testing	December	1992
International Standardized Profiles	January	1993

Preliminary Call for Papers

Interop Company is soliciting technical papers for an "Engineer's Conference" to be held as part of the upcoming NetWorld®+ INTEROP® 94 event, May 2-6 1994 in Las Vegas. The Engineer's Conference, which will run May 2nd and 3rd, is a two-day event offering solutions to practical systems/software design aspects of networking. All participants in the conference will be able to attend the NetWorld+ INTEROP 94 exhibition, which runs May 4-6.

Format

The conference will feature the presentation of original papers which will have been selected by a review committee. All accepted papers will be published in Conference Proceedings. Accepted papers have to be presented by original authors during the 2-day conference. The Engineer's Conference will focus on solutions to engineering design problems in three areas: High Speed Networking, Internetworking, and Network Management. This conference is making an effort to bring together research scholars, engineers, and vendors to address the engineering issues in the field of networking. It is an excellent forum for Engineers and Researchers to publish papers on solutions to today's engineering-related problems.

Topics

Papers are solicited in the following areas:

- *High Speed Networking*: ATM, Fast Ethernet, SONET, FDDI-II, HIPPI, SMDS, Frame Relay, Broadband ISDN, etc.
- *Internetworking*: Design of Bridges, Routers, and Multiprotocol Brouters, Addressing Schemes, Routing Protocols, Application Gateways etc.
- *Network Management*: Bandwidth utilization, Traffic Trend Analysis/Capacity Planning, Automated Trouble Ticket Systems, Congestion detection, Network Simulation, SNMP v1 and v2, Security, Export considerations for secure systems, Manager-to-manager communications, Standardized Testing Suites, Expert Systems, Accounting, Distributed/Hierarchical Management architectures, etc.

Submission guidelines

Interested authors are invited to submit an abstract (up to 100 words) clearly describing the problem and the solution offered. All abstracts will be reviewed and authors are notified for acceptance or rejection of the abstract. Authors of accepted abstracts must submit the paper before the last date. These papers are reviewed by a technical committee for technical merit of the paper before final acceptance.

Please note the important dates for abstract and paper submission. All abstracts *must* contain the authors name, address, telephone number, Fax number and e-mail address (if available). Please send your abstract to:

Interop Company
480 San Antonio Road
Mountain View,
CA 94040-1219 USA
Attention: Engineer's Conference
or e-mail it (in ASCII or *PostScript*) to: engineer@interop.com

Important dates

Abstracts due:	October 15, 1993
Notification to authors:	November 15, 1993
Draft paper due:	January 15, 1994
Feedback to authors:	February 1, 1994
Camera ready copy due:	February 20, 1994
Overhead slides due:	March 15, 1994

Authentication Code Available

Introduction

In the January 1993 issue of *ConneXions*, I described the hardware and software infrastructure that we use to support SLIP at the University of Minnesota. One of the major software parts is the “fanout” program that accepts authentication requests and forwards them to the correct hosts for verification. These requests are originated by Cisco terminal servers or by regular hosts for other purposes.

Fanout software

This “fanout” software has now been made available. In a nutshell, the fanout program accepts requests, either via the UDP-based TACACS protocol or via a TCP variant of that protocol. It performs some basic verification on the request and, if acceptable, forwards the request to a target host for yes/no verification. This design allows us to operate a University-wide service without having to have a single list of accounts and passwords.

Mechanism

In more detail, requests come from sources called “pools” (as in modem pools). To improve security, only requests from a specified list of sources will be accepted. These modem pools are also used as the basis of statistics gathering: there are utility programs to compute and graph pool usage.

The requests are validated in a number of ways. First, the request must be syntactically valid. Next, the user must not be on a list of “blocked users.” Next, the user must not be “anonymous,” “ftp,” “guest,” or “root.” Finally, the password must be non-null.

Protocols supported

The username is specified as an e-mail address: user@host. The “host” must be on a specified list of hosts. Also, the modem pool that the request comes from must be acceptable to that host. (Thus, the fanout supports restricted modem pools for, say, 800 service.) Finally, after all these checks, the user name and password are passed off to the host for verification. This pass-off can use any of these protocols:

- UDP-based TACACS
- FTP
- PopMail Version 2
- PopMail Version 3

The central fanout handles all access control checking, logging, and intrusion-tracking functions.

We also use this same fanout for other types of validation. For example, we have the ability to validate whether a person is a student or staff member of one of our campuses. Additional validation types can be added as necessary.

Included with the fanout program are a *ping* program for testing, a program that scans and digests the log files, and a program to compute and graph the modem pool usage.

Documentation

I am currently working on an informational RFC that documents the TACACS protocol. A copy of the draft is included in the distribution.

As described in the original article, the system does assume that you have a reliable, trustworthy IP routing substrate to work with. The additional verification also assumes a reliable electronic directory.

Availability This software is written in *Perl* and available for anonymous FTP from mail.unet.umn.edu in:

export/tacfanout.tar
export/tacfanout.tar.z (gzip'ed)

Mailing list There is also a mailing list for issues relating to the software:

disc-tacfanout@mail.unet.umn.edu

Please send requests to be added to the usual -request form.

Craig A. Finseth
Networking Services
Computer and Information Services
University of Minnesota
130 Lind Hall, 207 Church St SE
Minneapolis MN 55455-0134
USA
Desk: +1 612 624 3375
Problems: +1 612 625 0006
Fax: +1 612 626 1002
E-mail: fin@unet.umn.edu

Book Review

TCP/IP—Running A Successful Network, by K. Washburn and J. T. Evans, Addison-Wesley 1993, ISBN 0 201 62765, 459 pages.

This book is a recent addition to the many volumes on TCP/IP appearing over the last few years. It is written by two practitioners/implementors with good experience. It is largely of information in a readable form and seems to be very accurate. This is a good alternative text to the Comer Tri-Tome.

Organisation It is in 19 chapters divided into two halves:

The first half comprises 8 chapters and describes architectural principles, networking, protocols, addressing, routing and upper layer services. This section is oak, but i think covered better in other more general texts, but that is not a major criticism—you could use this as you single introductory textbook for students or new engineering/comms. software staff perfectly reasonably.

The second half is technology oriented and covers in the remaining 10 chapters, the details of the protocols and mechanisms from IP mapping onto LAN, WAN, ISDN etc., IP, TCP, Applications, DNS, NFS, RIP and EGP and OSPF and others, SNMP, Configuration and the Future. Each of these chapters ends with a list of relevant RFCs.

10 Appendices contain information on Contacting NICs, getting RFCs, subnetting procedures, Standards and some extremely useful and educational protocol traces.

Minor nit: For a book written in 1992/1993, the section on the future could contain mention of SIP/PIP/Nimrod etc...

Clear and readable Major plus: Clear exposition from theory to practice.

—Jon Crowcroft, University College London

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute



Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

480 San Antonio Road, Suite 100

Mountain View, CA 94040 U.S.A.

415-941-3399 FAX: 415-949-1779

connexions@interop.com

Back issues available upon request \$15./each
Volume discounts available upon request

CONNEXIONS